Mixed-Initiative Interaction for Maintenance Support

Amy Aragones [1], Joseph Price[1], Andrew Crapo[1], Nicholas Roddy[2] and Anil Varma[1]

[1]GE Global Research
1 Research Circle, Niskayuna, NY 12301
{aaragones,pricejo,crapo,varma}@research.ge.com

[2] GE Transportation Systems
East Lake Road 12-2 Erie, PA 16531
nicholas.roddy@trans.ge.com

Abstract. Automated maintenance support is a promising area for building and testing mixed initiative interaction systems. Today's maintenance decisions are increasingly proactive in nature, taking constantly evolving information into consideration to try to delay or avoid impending failure. A wealth of domain knowledge as well as situational awareness is required to make the right maintenance decisions in such a dynamic environment. We describe our work to date on the 'Reachback' system – a maintenance decision support tool that incorporates ideas from Case-Based Reasoning, decision theory, and adaptive interfaces in shaping its interaction with the user.

1 Introduction

Mixed-initiative interaction designs attempt to balance the automation of user needs and maintenance of user control over actions and decisions. These can be described as multi-agent systems that enable teams of agents (which may or may not include human agents) to collaboratively execute intelligent actions. Each agent contributes to the overall objective by performing the tasks that it does best and by taking initiative at the appropriate time. Mixed initiative interaction is contrasted against single initiative systems, which specify in advance which agent has the initiative in interaction. Horvitz (1999) refers to the mixed initiative approach in the design of human-computer interfaces as a "creative integration of direct manipulation and automated services… characterized by deeper, more natural collaborations between users and computers." This paper focuses on the scenario where an interface agent collaborates with a human agent in constantly evaluating an operational piece of equipment to determine if a maintenance intervention is necessary, its objectives achievable and its consequences positive. The basis for driving mixed initiative behavior is decision theoretic in nature, similar to that of Druzdzel (1999).

GE Global Research is engaged in developing technology to support the operations of its 13 businesses, varying from GE Aircraft Engines to the NBC television network. A constant and significant area of attention for the past few years has been the maintenance domain, so that industrial equipment varying from MR scanners to locomotives, power turbines and aircraft engines can be serviced efficiently and in a cost-effective manner. Advances in information digitization and storage and accessibility of high bandwidth communication have vastly increased the quantity of information that is available regarding a given maintenance event. The nature of maintenance intervention has also changed. From focusing on reactive response – fixing a piece of equipment after it breaks – the emphasis has expanded to include predictive monitoring and proactive repair.

In such a scenario, a maintenance decision maker has to evaluate the impact of constantly updated operational information in deciding if the time is right for a maintenance intervention or if the equipment should be allowed to continue operating until additional evidence or confirmation of imminent failure is presented. Now consider this task replicated across hundreds of such pieces of monitored equipment, each of which could be operating normally, experiencing transient problems or be potentially in different stages of component or system failure, as reported by an imperfect monitoring system. The need for a mixed initiative system arises here, as in many other cases, in keeping the decision-maker focused on the few issues that can best impact his or her objectives. Assuming that this mediation is performed by an 'interface agent', the task of the agent can be outlined as selecting the most appropriate information and display format at the right time to support the human's decision process, and offering decision strategies generated by an underlying reasoning component that offers the best outcomes.

2 Application Domain

The work described in this paper supports a maintenance concept called "Reachback". The term "Reachback" has its origins in military operations. Neal (2000) defines Reachback as "the electronic ability to exploit organic and nonorganic resources, capabilities and expertise, which by design are not located in theater." Making better decisions through access to and assimilation of remote, asynchronous sources of information is a defining feature of a reachback capability. Our work seeks to apply reachback concepts to "Service and Support" - the task of maintaining deployed equipment while balancing competing constraints of low maintenance cost and high equipment availability.

The initial prototype of the Reachback system discussed here uses data from GE Transportation Systems. GE Transportation currently monitors over 4000 locomotives around the world. As faults and other anomalous conditions develop onboard the locomotive, they are collected and downloaded to a central monitoring center for analysis. An engineer examines this data and the output of multiple AI based diagnostic tools to diagnose if the equipment is experiencing a problem. If a problem is indeed identified, a recommendation is made to have a field engineer perform troubleshooting and maintenance on the locomotive.

The next generation of military hardware being designed today is expected to have similar capabilities for monitoring and information communication. Though this represents a vast increase in the amount of visibility into the equipment's condition, it is not always obvious what the right maintenance decision is. The sensors on the equipment usually indicate 'anomalies' rather than hard failures. In addition, environmental and operating conditions have to be taken into account when distinguishing actual failures from false alarms. The importance of the mission affects whether it is practical to take the time to perform maintenance, as does the availability of personnel and parts.

There are two primary audiences for a maintenance reachback system. The first is either the user who either operates the equipment directly or whoever is responsible for remotely monitoring it. Either of these have the authority to order maintenance to be performed, and are responsible for making informed decisions regarding the scope and timing of the maintenance. Once the decision to perform maintenance has been made and the equipment is in the repair facility, then a field engineer (the second user) can use guidance in performing the best mix of reactive and preventive maintenance so that equipment availability/cost of ownership is optimized.

We believe that intelligent maintenance support represents a natural domain for mixed-initiative systems. A reachback system relies on conveying a very large information set about virtual assets – much of which may be uncertain – to the decision maker in a manner that is conducive to making the best decision. The advantages of the reachback system derive from making a larger quantum of information exchange happen as compared to traditional logistics systems. However, this potentially poses the problem of providing too much information for the human user to efficiently access and assimilate. The nature of information overload is fundamentally dynamic since the decision process, and therefore the need for decision support, changes in response to the context in which decisions are being made. Experts in a particular subject matter, for example, are likely to have learned how to navigate and filter the increasingly large amount of diagnostic data reported by a monitoring and diagnostic system, procedural information and other supporting material. Novices, on the other hand, lack this experience and are more likely to need simpler data presentations but different, potentially additional supporting documentation and learning materials. Information needs may also change according to the confidence users have in the accuracy of the data presented, which may in part be driven by their level of expertise as well as the confidence the system itself has in the data. In addition, what is likely to be an appropriate level of information for a service visit in a shop may not be an appropriate level of information for service while the equipment is in the field, where time is limited.

In part, this problem can be addressed through intelligent fusion and bundling of information through a reasoning system deeper within the reachback system. However, the manner in which these information bundles are conveyed to the user is also critically important since, particularly in open-ended work scenarios, it is the human user who will be making the decisions about what to do with the information. Interaction with the system may include requests for more information in order to respond to variables that the reachback system cannot observe, to clarify why the system reached particular conclusions, to establish confidence in the information being delivered, or to visualize information in different ways to enable cognitive processing. However, because reachback supports knowledge-based work across a varied user population in multiple domains where new scenarios develop over time, it becomes risky to assume that we can define information requirements – not only what information to deliver, but also how and when – upfront in the design process.

There are several factors that are important to consider in modeling the interaction with the user:

- The nature of the information to be conveyed
- The expertise of the end user
- The preferences of the user, which may be based on sensory or cognitive capacities
- The time available to resolve a maintenance issue
- Confidence in the accuracy of the information delivered by the system

In particular, the user's expertise and confidence in decision support information are dynamic, variable and, to some degree, directly unobservable by the system. Adaptation to these contextual constraints then, have been identified as key challenges in the interactive process between humans and computers in a reachback system.

Horvitz (1999) outlines certain critical factors that favor the use of mixed initiative user interfaces. We believe that a good way to describe our maintenance domain is to define its characteristics against some of these principles laid out by Horvitz.

1. **Developing significant value-added automation:**

At the very least, an intelligent maintenance system should automate several routine checks and calculations that verify that the equipment being monitored does indeed require attention. Examples would include checking that an alert has not already been issued for the equipment, or that the equipment is not operating in unseasonable temperatures that can cause failure-like indications.

2. **Considering uncertainty about a user's goals:**

To some extent, a user's goals can be inferred from his/her role definition encoded in the system. Understanding if the user is still in the 'diagnosis' phase and trying to isolate the problem, or looking to see if the right parts can be obtained can be key inputs into how the maintenance system takes initiative. Users may also jump back and forth between different hypotheses about the equipment condition. Providing related information to reduce uncertainty can be a key attribute of the reachback maintenance system.

3. **Considering the status of a user's attention in the timing of services:**

Time is a frequent constraint in maintenance actions. Being able to display information and guidance that degrades the quality of work gracefully with the time available is vital. In the case of maintenance, the status of a user's training in being able to assimilate complex technical instructions is also a factor. We refer to this capability as the adaptive aspects of the mixed initiative interface. Thus not only the sharing of control is important in mixed initiative, but also what Horwitz refers to as "scoping the precision of service" to match the user's goals.

4. **Inferring ideal action in light of costs, benefits, and uncertainties:**

The tradeoff between cost and benefit is at the heart of the maintenance problem formulation. The loss incurred in delaying maintenance after a machine has definitely failed can be computed with relative ease, since the time after failure is unambiguously known. As we move toward proactive and predictive maintenance, the cost of performing unnecessary maintenance and giving up available life is hard to compute, without a rigorous control sample that is allowed to go to failure under similar conditions. In the case of aircraft engines or power turbines, such control samples seldom exist.

Thus it is vital that all available information – whether design limits, experimental results, field failure data and impact on metrics be well modeled with as accurate an assessment of outcome probabilities as possible. This will ensure that any decision-theoretic calculation of the ideal action is also reasonable in real life.

5. **Maintaining a working memory of recent interactions:**

Any scalable system with multiple roles and users must have a component of learning that aids in refining the set of parameters that drive its behavior. We propose to use CBR to capture successful interactions with our Reachback system. As the system gains experience with multiple users and interaction scenarios, it will be better able to decide its limits for initiating interaction, as well as the format of the interaction.

3 The Reachback System

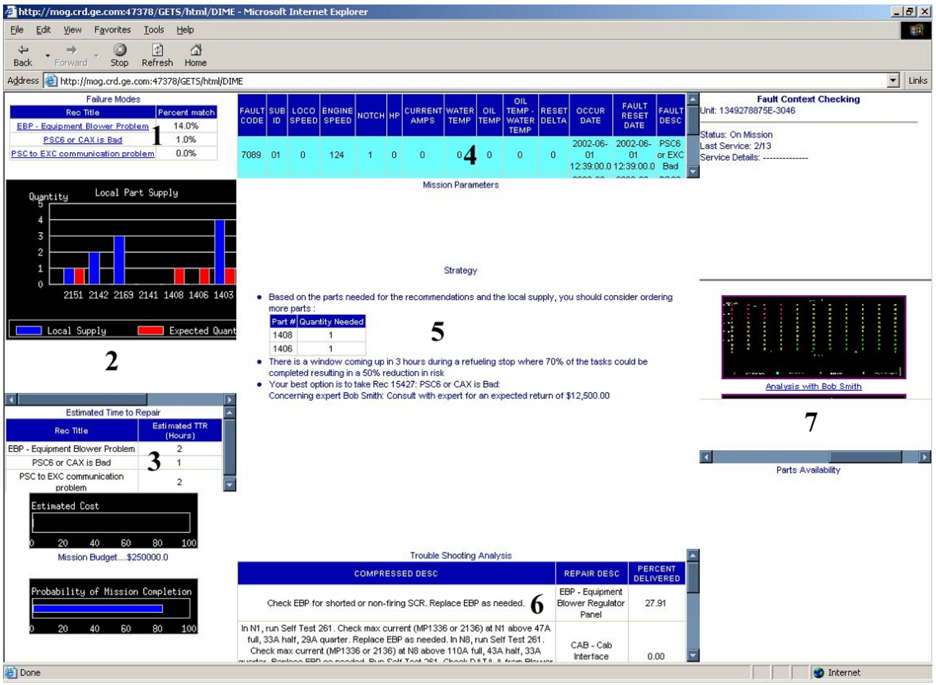

Figure 1: Reachback User Interface Screen Shot Showing Areas

This section describes the actual interface of the Reachback system. The system currently links to data on several hundred locomotives. Once the user picks a particular piece of equipment to monitor, he is given an assessment of the potential problems happening on that equipment, and an analysis of his/her options regarding whether to wait, order parts, order maintenance, consult an expert, or do a combination of all or some of these. A sample screen from the Reachback interface is shown in Figure 1.

a) Faults

FAULT CODE	SUB ID	LOCO SPEED	ENGINE SPEED	NOTCH	HP	CURRENT AMPS	WATER TEMP	OIL TEMP	OIL TEMP - WATER TEMP	RESET DELTA	OCCUR DATE	FAULT RESET DATE	FAULT DESC
7089	01	0	124	1	0	0	0	0	0	0	2002-06-01 12:39:00.0	2002-06-01 12:39:00.0	PSC6 or EXC Bad

Faults represent the primary indication that a problem may be occurring on the locomotive. Faults stream intermittently into the system from the different locomotives being monitored. The fault indicators are shown in area 4 of Figure 1.

b) Failure Modes

Failure Modes

Rec Title	Percent match
EBP - Equipment Blower Problem	14.0%
PSC6 or CAX is Bad	1.0%

The faults trigger analysis by a rule based and a case-based diagnostic system and a determination is made of possible failure modes and their likelihood. This information is displayed in area 1.

c) Part Supply

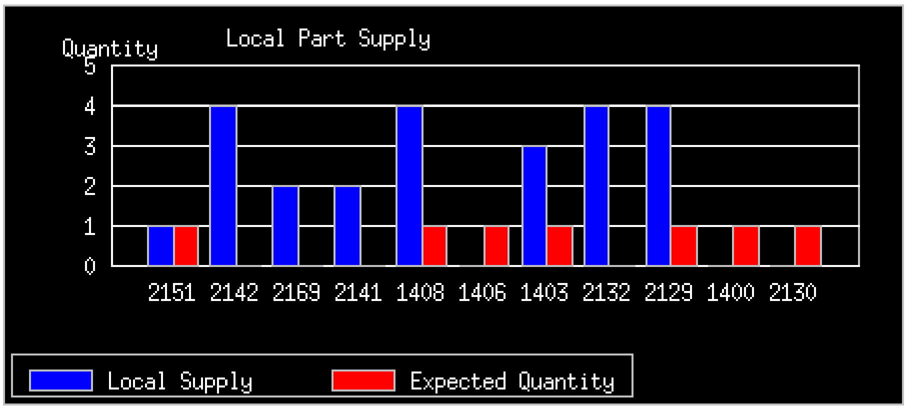

Based on expected failure modes, a graph of the expected quantity of parts needed is calculated. This expected quantity is evaluated against information about available stocking levels. If it appears that certain parts are likely to be in short supply, a supply agent initiates a search for needed parts among a population of remote suppliers. This is shown in area 2 in Figure 1.

d) Estimated Time

Estimated Time to Repair

Rec Title	Estimated TTR (Hours)
EBP - Equipment Blower Problem	2
PSC6 or CAX is Bad	1
PSC to EXC communication problem	2

Field validated times required to address the different failure modes are retrieved and are available as inputs for analysis. This is shown in area 3 in Figure 1.

e) Trouble Shooting

Trouble Shooting Analysis

COMPRESSED DESC	REPAIR DESC	PERCENT DELIVERED
Check EBP for shorted or non-firing SCR. Replace EBP as needed.	EBP - Equipment Blower Regulator Panel	27.91
In N1, run Self Test 261. Check max current (MP1336 or 2136) at N1 above 47A full, 33A half, 29A quarter. Replace EBP as needed. In N8, run Self Test 261. Check max current (MP1336 or 2136) at N8 above 110A full, 43A half, 33A	CAB - Cab Interface	0.00

For each failure mode, there is a set of troubleshooting procedures that must be performed on-site. These can help to further isolate the problem and more reliably identify the parts that may be required to fix it. This is shown in area 6 in Figure 1.

f) Expert Analysis

14.0% 5603 EBP - Equipment Blower Problem 2
1.0% 15427 PSC6 or CAX is Bad 1
0.0% 8501 PSC to EXC communication problem 2

Bob Smith 95%
Susan Jones 80%
Larry Thompson 60%

analyze scenario

The expert analysis module keeps an account of domain experts in specific areas of troubleshooting, i.e., engine experts, propulsion experts, etc. and their historical track record in making successful diagnoses. This module performs a Value of Perfect Information analysis to determine if consulting an expert to validate the output of the diagnostic tools is likely to benefit the decision process. The cost of parts is traded off vs. the cost/likelihood of an equipment failure if maintenance is not performed. This is shown in area 7 in Figure 1.

g) Strategy

Strategy

- Based on the parts needed for the recommendations and the local supply, you should consider ordering more parts :

Part #	Quantity Needed
1408	1
2151	1

- There is a window coming up in 3 hours during a refueling stop where 70% of the tasks could be completed resulting in a 50% reduction in risk
- Your best option is to take Rec 15427: PSC6 or CAX is Bad:
  Concerning expert Bob Smith: Consult with expert for an expected return of $12,500.00

Finally, the system assembles a strategy that looks like the following. This is initially based entirely on the information available to it without any input from the user. This is shown in area 5 in Figure 1.

4 Mixed Initiative Interaction and Case Based Reasoning

In practice, the Reachback tool is expected to monitor a number of events happening in the background. Examples of such events would include

- A new fault occurs, re-executing the diagnostic engine with this new input, thereby confirming an existing problem or indicating the presence of a new one. This initiates a number of tasks for the parts supply and expert analysis agents, among others.
- The knowledge within the diagnostic case-base gets updated. This may increase or decrease the confidence of the system in a newly diagnosed failure mode.
- An expert might come on line. Depending on their specialty, it may make sense to initiate a collaboration.
- Another user has spent some of the shared resource granted for this mission. This restricts new options available to the current user. For example, someone else has ordered parts for a different problem.
- The local supply of a part might change. The system may or may not need to alter strategy based on this update.

The challenge is to incorporate this new information into the decision to initiate an interaction, or to modify an interaction in progress. Our preferred approach for maintaining and evolving this system is to have users interact with it and capture their sessions. A set of archived “interactions” – or cases – constitute the knowledge base from which the system learns how to best respond to new situations. Attributes of a case include

a) Static information: User, Role, Equipment type etc.
b) Preference for visual layout, nature of information displays.
c) Mission specific parameters and constraints.
d) Patterns of exploration for additional information.
e) Successful or unsuccessful outcomes, if available.

Underneath the Reachback system lies a maintenance ontology that defines the entities and relationships within the system. This ontology combines the top-down approach of Sowa (2000) with an integrated domain ontology. Of particular importance in the human-computer interface is a model of the communication process. Communication is seen as a process in which an agent (the initiator) encodes information in a semiotic sign or text and transfers this to a second agent (the recipient) for decoding and interpretation (Sowa 2000, Chandler 1994). The human user and the Reachback system represent the two participants of the mixed interaction communication, with information flow occurring in either direction. The “sign type” is a subclass of “schema,” whose subtypes include “visual signs.” “Visual signs,” in turn, can be “graphical,” “tabular,” or “textual.” The model of communication demonstrates the basic ontological concepts of taxonomic hierarchy and other kinds of association. An ontology, especially with inheritance, allows the system to reason at a more abstract or a more detailed level to best match the information available. For example, if no data exists on the kind of information display a particular user prefers, the system can explore available information on "similar" users-- users belonging to the same ontological classification-- to generate its response.

5 Conclusions

There is a wealth of raw and processed information available to the maintenance decision maker, but it is of little value unless its availability results in better quality decisions leading to measurably improved operations. A mixed initiative system that carefully considers the value of existing and incoming information, as well as knowledge from prior interactions, to customize its response to its users is a promising approach. We have described the ongoing development of the Reachback system for maintenance support that is designed to use prior cases to drive its mixed interaction model and settings.

6 Acknowledgements

This work is supported by Air Force Research Labs and by GE Transportation Systems.

References

1. Chandler, Daniel. Semiotics for Beginners [WWW document] URL http://www.aber.ac.uk/media/Documents/S4B/ [August 7, 2002] (1994).
2. Druzdzel, Marek J. "ESP: A mixed initiative decision-theoretic decision modeling system," In Working Notes of the AAAI-99 Workshop on Mixed-Initiative Intelligence, Orlando, FL (18 July 1999) 99-106.
3. Horwitz, E. "Principles of Mixed-Initiative User Interfaces," Proceedings of CHI '99, ACM SIG-CHI Conference on Human Factors in Computing Systems (1999) 159-166.
4. Neal, John M. "A Look at Reachback," Military Review, Volume LXXX, Number 5, September-October (2000).
5. Sowa, John F. Knowledge Representation: Logical, Philosophical, and Computational Foundations, Brooks/Cole, Pacific Grove, CA (2000).